

# On the Predictability of Talk Attendance at Academic Conferences

Christoph Scholz   Jens Illig   Martin Atzmueller   Gerd Stumme

Knowledge and Data Engineering Group, University of Kassel  
Wilhelmshöher Allee 73, D-34121 Kassel, Germany  
{scholz, illig, atzmueller, stumme}@cs.uni-kassel.de

## ABSTRACT

This paper focuses on the prediction of real-world talk attendances at academic conferences with respect to different influence factors. We study the predictability of talk attendances using real-world tracked face-to-face contacts. Furthermore, we investigate and discuss the predictive power of user interests extracted from the users' previous publications. We apply *Hybrid Rooted PageRank*, a state-of-the-art unsupervised machine learning method that combines information from different sources. Using this method, we analyze and discuss the predictive power of contact and interest networks separately and in combination. We find that contact and similarity networks achieve comparable results, and that combinations of different networks can only to a limited extent help to improve the prediction quality. For our experiments, we analyze the predictability of talk attendance at the ACM Conference on Hypertext and Hypermedia 2011 collected using the conference management system CONFERATOR.

## 1. INTRODUCTION

Academic conferences facilitate scientific exchange, collaboration and innovation, e.g., fostered by social contacts and interesting talks. A major task for every conference attendee is the selection of talks relevant to his research. Conference guidance systems such as *Conference Navigator* [28] and CONFERATOR<sup>1</sup> [3], support this with the possibility of creating a personalized schedule. Picking talks manually, however, may become complex due to the large amount of available talks at a conference. Furthermore, conversations with other attendees and changes in the conference schedule can influence the talk selection.

Recommendation components of conference guidance systems can support their users by presenting suggestions of talks which the system determined as most interesting for the respective user. Then, such recommendations influence the decision e.g., due to recommended talks which were otherwise not considered by the user. Therefore, recommender systems should ideally always be evaluated in an online scenario, where influence is part of the evaluation.

<sup>1</sup><http://www.conferator.org>

In this paper, we focus on the predictability of real talk attendances, i.e., we try to find models imitating the actual decision process without recommendation influence. Due to the low availability of online recommender evaluation, it is reasonable to evaluate recommender systems on a prediction setting. This is partially valid since good predictions are also good recommendations to the extent that the user does not regret the predicted decisions. For our evaluation, we use real-world talk attendance data which was collected using CONFERATOR. CONFERATOR applies active RFID technology developed by the SocioPatterns consortium<sup>2</sup> for the localization as well as for the measurement of face-to-face contacts between researchers during the conference, e.g., during the coffee breaks. Based on such RFID data and collected content information of scientific papers, we investigate the potential of social contact information and content-similarity for predicting real-world talk attendance decisions. Especially, we analyze the potential of combining different information sources for improving the overall prediction quality.

Our contribution can be summarized as follows:

1. We present the first study about the predictability of visited talks at academic conferences on real world data.
2. We analyze different influence factors concerning the predictability of talks at academic conferences. In particular, we study the influence of face-to-face contacts and user interest on the talk attendance decision.
3. We consider and adapt state-of-the-art unsupervised link prediction methods for the talk prediction problem, focusing on the *rooted PageRank* [18] and the *Hybrid Rooted PageRank* [9] algorithms.
4. We present an in-depth analysis of talk attendance predictability using different performance metrics and investigate the influence of different interaction networks, e.g., derived from social contact and content information, for this task.

The rest of this paper is structured as follows: Section 2 discusses related work. In Section 3 we describe the framework that we used to collect our data. Section 4 gives a detailed overview of the collected dataset. In Section 5, we discuss the algorithms used for the prediction task. After that, Section 6 presents a detailed evaluation using the dataset collected at ACM Hypertext 2011. Finally, Section 7 summarizes our results.

<sup>2</sup><http://www.sociopatterns.org>

## 2. RELATED WORK

In this section, we discuss related work concerning the talk prediction problem at academic conferences. We start with relevant work about the analysis of human contact pattern at conferences and then discuss work about talk recommendation.

### 2.1 Analysis of Human Contact Patterns and Link Prediction

The analysis of offline social networks, focusing on human contacts, has been largely neglected. In this context, Eagle et al. [11] and Zhoe et al. [13] presented an analysis of proximity information collected by devices based on Bluetooth communication, similar to Xu et al. [29], who also related this to online social networks. However, in all these experiments it was not possible to detect reliable face-to-face contacts. The SocioPatterns collaboration developed an infrastructure that detects close-range and face-to-face proximity (1-1.5 meters) of individuals wearing proximity tags with a temporal resolution of 20 seconds [8]. Due to the fact that the human body blocks RFID signals this allows the detection of face-to-face contacts between persons. One of the first experiments using this kind of proximity tags was done by Catuto and colleagues in [2]. They presented an application that combines online and offline data from conference attendees. In [30] the authors also studied the influence between offline and online properties using a mobile social application in the context of academic conferences. Barrat et al. [6] compared the attendees' contact patterns with their research seniority, their co-authorship and their activity in social web platforms. The SocioPatterns sensing infrastructure was also deployed in other environments in order to study the dynamics of human contacts, such as healthcare environments [14], schools [26] and museums [15]. Atzmueller et al. [4] described the dynamics of community structures and roles at conferences, extending the analysis of interactions and dynamics, and the connection between research interests, roles and academic jobs of conference attendees [19].

Link prediction, as defined by Liben-Nowell and Kleinberg in [18], is strongly related to talk prediction. In [18] the authors did a first comprehensive analysis by analyzing the predictability of unsupervised machine learning methods. Scholz et al. analysed the predictability of face-to-face contacts at academic conferences. In [9] the authors presented an unsupervised link prediction method that combines information of different networks.

### 2.2 Talk Recommendation and Prediction

To the best of our knowledge, predictability of scientific talk attendance has not yet been investigated w.r.t the true physical attendance of conference talks. Talk recommendation is a specific instance of the general recommendation task. Published work about talk recommendation systems evaluated the recommendation algorithms with respect to their ability to reconstruct the remaining part of a partly given user's attendance plan entered into some conference management system.

For recommender systems, we typically distinguish between content-based and collaborative-filtering approaches [1]. Content-based recommenders make use of properties of the recommended items, while collaborative filtering methods utilize common item ratings of users. For talk recommendation, items are talks, while author, title, and abstract are content-properties.

Minkov et al. [21] as well as Pham et al. [22] simulated talk at-

tendances and collected explicit user feedback in form of questionnaires about the generated recommendations. Based on this feedback they evaluate their algorithms. Both evaluation schemes have their drawbacks. Using attendance plans as a gold-standard is not absolutely correct, because it is unclear whether the user actually attended the talks. Instead, plans may be non-final or simply be a collection of bookmarks used as a reminder for later attendance decisions. Furthermore, the plans may be incomplete in terms of not covering all time slots of the conference. Using questionnaires, the user is usually only asked once or few times about her satisfaction with the recommendations. Therefore, questionnaires give only a very rough measure about the overall recommendation quality, contrasting the evaluation for the prediction task where there are quantitative evaluations based on the correctness of each single prediction. In questionnaires, users can also rate the recommendation quality high if the system recommended a talk which the user did not attend but still found interesting. The other way round, attended talks may have bored the user and could therefore be bad recommendations although being perfect predictions. Questionnaires therefore measure a slightly different property than measured when evaluating in a prediction setting.

Minkov et al. [21] trained a RankSVM [16] classifier by supervision from a training part of their user feedback and evaluated on a test set. They augmented their content-based approach with a collaborative aspect using a modified RankSVM optimization problem which integrates dimensionality reduction and optimizes the dimensionality reduction parameters across users. In contrast to this work, our work uses a more explicit usage of social networks and focuses on unsupervised or weakly supervised approaches since we expect little knowledge about the talk attendance preferences of most conference visitors.

Pham et al. [22], as well as Lee and Brusilovsky [17] applied collaborative filtering for the recommendation of talks. For each user, [22] calculate sets of the 5 most similar users according to either commonly bookmarked talks in *Conference Navigator* or common co-authors. They furthermore use content-boosted collaborative filtering. User-similarity is then not solely calculated from known shared bookmarks or known co-authored publications (with a weight of 1), but also from relations to other users which have bookmarks of talks or co-authored papers similar to those the target user is related to (using the maximum cosine similarity to any of the other users bookmarks). Results were further enhanced by reweighting scores based on numbers of common co-authors and co-authored papers but precision was reported to be low with a maximum of 21%. Lee and Brusilovsky [17] rely on bookmarks obtained from the conference management system *Conference Navigator 2.0* [28]. This approach also uses boosted collaborative filtering. But instead of adding content similarity to the similarity based on bookmarks (or planned attendances), [17] calculate most similar users based on a weighted average of Jaccard coefficients on common co-authors and commonly referenced publications. They evaluated recommendation quality based on *conference simulations* where feedback from six *evaluators* was retrieved using evaluation forms. They provided two textual example statements but no quantitative evaluation.

The difference between our work and existing literature is that we present the first analysis of the predictability of visited talks at conferences using real world data. Furthermore, we study the influence of face-to-face contacts and user interest concerning the talk prediction problem. In particular, we consider combinations of different knowledge sources given as social interaction networks.

### 3. CONFERATOR – A SOCIAL CONFERENCE MANAGEMENT SYSTEM

In the following section, we first outline the active RFID technology used by the CONFERATOR system. Next, we introduce the CONFERATOR and its functionality.

#### 3.1 RFID-Setup

At the Hypertext 2011 conference we asked each participant to wear an active RFID tag (see Figure 1). One decisive factor of these active RFID tags is the possibility to detect other active RFID tags within a range of up to 1.5 meters, which allows us to create human face to face contact networks. We call these active RFID tags *proximity tags* in the following. Each proximity tag sends out two types of RFID-signals, proximity signals and tracking signals. A proximity signal is used for contact sensing, which is achieved by using signals with very low radio power levels [5]. The proximity tag sends out tracking signals in four different signals strengths (-18dbm, -12dbm, -6dbm, 0dbm) to RFID readers (see Figure 1) placed at fixed positions in the conference area. These tracking signals are used to transmit proximity information to a central server and for determining the position of each conference participant [24] [20]. Depending on the signal strength the range of a tracking signal inside a building is up to 25 meters. Each signal contains the signals strength and ID of the reporting tag and the IDs of all RFID tags in proximity. For more information about the proximity tags we refer to Barrat et al [5] and the OpenBeacon website.<sup>3</sup>



Figure 1: Proximity Tag (left) and RFID Reader (right)

#### 3.2 Conferator

The CONFERATOR system [3] is a social and ubiquitous conference guidance system. It allows conference participants to manage their conference schedule. Furthermore, the CONFERATOR supports social interaction at a conference. For example, it is possible for conference participants to recall their own contacts or to browse through other conference attendees' user profiles. Furthermore, the CONFERATOR presents personalized suggestions for interesting talks (see Figure 2). We note here that the recommendation component presented in Figure 2 was not a part of the CONFERATOR at HT 2011. Here the talk recommendations were provided by the Conference Navigator [28]. CONFERATOR has successfully been deployed at several events, e.g., the LWA 2010<sup>4</sup>, LWA 2011<sup>5</sup> and LWA 2012<sup>6</sup> conferences, the Hypertext 2011<sup>7</sup> conference, the INFORMATIK 2013<sup>8</sup> conference, and a technology day of the Venus<sup>9</sup> project. In this paper, we focus on data collected at the Hypertext 2011.

<sup>3</sup><http://www.openbeacon.org>

<sup>4</sup><http://www.kde.cs.uni-kassel.de/conf/lwa10/>

<sup>5</sup><http://lwa2011.cs.uni-magdeburg.de/>

<sup>6</sup><http://lwa2012.cs.tu-dortmund.de/>

<sup>7</sup><http://www.ht2011.org/>

<sup>8</sup><http://informatik2013.de/>

<sup>9</sup><http://www.iteg.uni-kassel.de/>

## Talk Recommender



Figure 2: Talk recommender in the social conference management system CONFERATOR.

### 4. DATASETS

In the following section we introduce the dataset collected at the 22<sup>nd</sup> ACM Conference on Hypertext and Hypermedia 2011 (HT 2011) in Eindhoven. We present statistics characterizing key properties of the applied data.

#### 4.1 Face-to-Face Contact Data

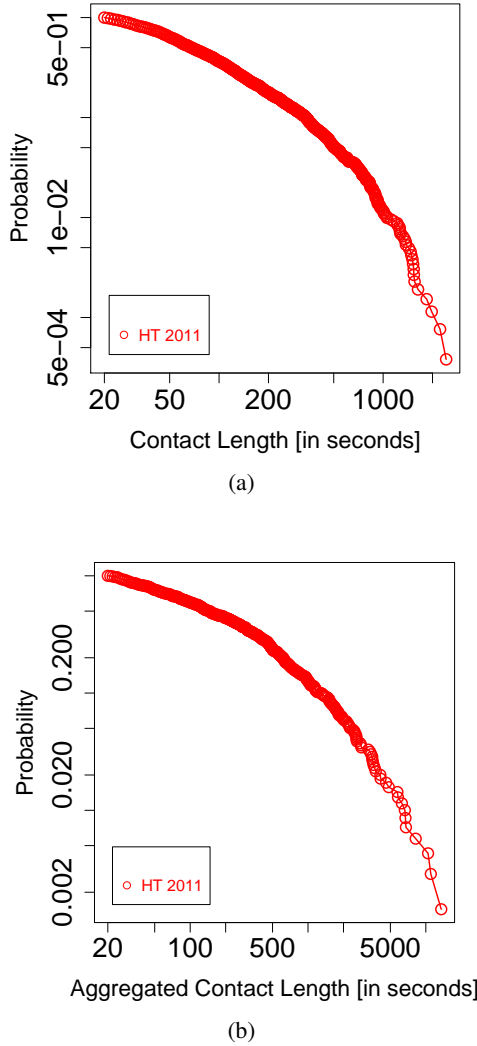
Table 1 provides a summary on the characteristics of the collected face-to-face proximity dataset. As already observed before [8, 15, 19], the distributions of all contacts and all aggregated face-to-face contacts lengths between conference participants are heavy-tailed (see Figures 3(a) and 3(b)). More than half of all aggregated face-to-face contacts are shorter than 200 seconds and the average contact duration is less than one minute. However, very long contacts are also observed. The diameter, average degree, and average path length of  $G$  are similar to the results presented in [4, 15]. For more details on the applied dataset, we refer to, e.g., [9] and [19].

	HT 2011
#days	3
$ V $	68
$ E $	698
Avg.Deg.( $G$ )	20.53
APL ( $G$ )	1.76
$d(G)$	4
AACD	529

Table 1: Collected dataset at HT 2011. Here  $d$  is the diameter, AACD the average aggregated contact-duration (in seconds) and APL the average path length.

#### 4.2 Talk Attendance Data

For our analysis, we focus on the parallel talks at HT 2011. Overall, 14 parallel talks took place in two rooms. For our prediction analysis it is essential to determine whether or not a participant attended a talk. Therefore, we installed one RFID reader in each conference room. As described in Section 3.1, a proximity tag sends out tracking signals that we used to determine the current position of each conference participant at room level basis. For the determination of talk attendance we used the following localization strategy: Since




**Figure 3: Cumulated contact length distribution of all face-to-face contacts (a) and all aggregated face-to-face contact at the HT-2011 conference. The  $x$ -axis displays the minimum length of a contact in seconds, the  $y$ -axis the number of contacts having at least this contact length, respectively. The axes are scaled logarithmically.**

the walls of each conference room (where the talks took place) were very thick and hence tracking signals (sent out by the proximity tags) could only be detected within one conference room. This means, that if we detect a tracking signal of a conference participant in a conference room, then we know that that this participant must be in this room. Overall, we observed 359 visited talks from 53 conference participants.

### 4.3 Full-Text Data

For our prediction task, we also consider the content of all papers. For each conference participant, we therefore crawled all papers that are listed in DBLP since 2006. In total, we crawled 707 papers. With the full-text data we created bag-of-words models representing the paper profiles for each participant. For the participants' bag-of-words-model construction, we used the Porter Stemmer al-

gorithm [27] and removed all stop words. Figure 4 displays the cumulative number of papers for each conference participant.



**Figure 4: Distribution of the number of papers for each conference participant. The  $y$ -axis displays the number of participants having at least the number of papers that are defined by the  $x$ -axis.**

## 5. ALGORITHMS

In this section, we describe the algorithms used for the prediction of talks at academic conferences. Focusing on unsupervised methods, we use the *Hybrid Rooted PageRank* algorithm, an extension of the *rooted PageRank* algorithm, for prediction.

### 5.1 Rooted PageRank

The *rooted PageRank* predictor (RPR) [18] is an adaption of the *PageRank* algorithm [7] for the link prediction task. The *rooted PageRank* predictor score between participants  $r$  and  $y$  is defined by the stationary probability distribution of participant  $y$  under the following random walk [18]:

- With probability  $\alpha$ , jump to  $r$ .
- With probability  $1 - \alpha$ , jump to a random neighbor of the current node.

For the *weighted rooted PageRank* (WRPR) predictor, the random walk selects a random neighbor  $n$  of the current node  $c$  with probability  $\frac{w(c,n)}{\sum_{c \rightarrow d} w(c,d)}$ , where  $w(c,d)$  is the weight of the edge  $(c,d)$ .

### 5.2 The Hybrid Rooted PageRank Method

In this section, we describe the *Hybrid Rooted PageRank* algorithm, first presented in [9]. This algorithm is an unsupervised machine learning method and extends the *rooted PageRank* algorithm. The *Hybrid Rooted PageRank* algorithm combines the information of different networks. To do so, the *Hybrid Rooted PageRank* computes the stationary distribution of nodes under the random walk described in Algorithm 1. In each step, the walk selects a given network with respect to a given probability distribution. From the current node  $c$  a link in this network is then selected to a random neighbor  $n$  of node  $c$  with probability  $\frac{w(c,n)}{\sum_{c \rightarrow d} w(c,d)}$ , where  $w(c,d)$  is the weight of the edge  $(c,d)$ . If no link exists in the chosen network (i.e., if the node is isolated), then the algorithm jumps back to

**Input** : Networks  $N = \{N_1, \dots, N_n\}$ ,  
Network-Probabilities  $P = \{p_1, \dots, p_n\}$ ,  
Probability  $\alpha$ , Root node  $r$

**Output**: Stationary distribution weight of node  $v$  under the following random walk:

- 1 With probability  $\alpha$  jump to root node  $r$ .
- 2 With probability  $1 - \alpha$ :
- 3 Choose Network  $N_i \in N$  with respect to probability distribution  $P$ .
- 4 **if** There exist no outgoing edges **then**
- 5     Jump to root node  $r$
- 6 **else** From the current node  $c$  jump to a neighbor  $n$  selected with a probability  $\frac{w(c,n)}{\sum_{c \rightarrow d} w(c,d)}$ , i. e., proportional to the weight  $w(c, n)$  of the edge  $(c, n)$ .

**Algorithm 1:** Hybrid Rooted Random Walk

the root node. In this way, one can integrate different networks for prediction of links.

Assume we want to determine the *Hybrid Rooted PageRank* predictor score for participants  $r$  and  $x$ . In this case, we use participant  $r$  as root node and execute the algorithm. As a result, the algorithm computes the stationary probability distribution of all nodes. The predictor score between participant  $r$  and  $x$  is then given by the stationary probability distribution of participant  $x$ .

## 6. EVALUATION

In this section, we analyze the predictability of talk attendance at academic conferences. Specifically we study the influence of face-to-face contacts and user interests on this prediction problem. Furthermore we consider combinations of different knowledge sources given as social interaction networks. We start with an explanation of the used evaluation methods, before we present and discuss the predictability results.

### 6.1 Evaluation Method

In this section we define and discuss the measures that we calculate for evaluating the impact of various examined influence factors on real world talk attendance decisions. We use two measures: *Accuracy* and *area under the curve* of the *receiver operating characteristic*.

#### 6.1.1 Accuracy

Accuracy (ACC) is widely used and simply refers to the fraction of correct decisions divided by the total amount of decisions.

The problem of predicting attended talks, which we cover in this work, is distributed over time slots. Naturally, for every person only one talk can be attended at each time slot. Thus, for each time slot, one decision *has* to be made for every conference participant that attends one of the parallel talks. Applied to our talk prediction setting, accuracy can be interpreted as the maximum likelihood probability estimate for a talk recommender system to correctly predict the next attended talk.

#### 6.1.2 Area Under the Curve of the Receiver Operating Characteristic

We further use the area under the receiver operating characteristic (here simply abbreviated as AUC) [12]. The receiver operating

characteristic (ROC) is given by a plot, which is defined as follows. For each  $n$ , a point is added to the curve, based on the top  $n$  decisions of the algorithm ranked by relevance. The  $x$ -coordinate of the point is the false-positive rate of these  $n$  decisions, and its  $y$ -coordinate is the true-positive rate.

For our talk prediction task, we use AUC to evaluate a conference-global ranking. For each pair of parallel talks  $t_1$  and  $t_2$  and talk attendee  $p$  we calculate two predictor scores, one for  $t_1$  and one for  $t_2$ . This results in a ranking containing all positive and negative decisions for predicting all talks. AUC evaluation rewards a predictor's ability to rank correct decisions before wrong decisions according to the ground truth. An ideal predictor ranks all correct decisions above all wrong predictions and achieves thus an AUC score of 1.0, while a purely random predictor achieves a score of 0.5.

## 6.2 Predictability of Talk Attendance

In this subsection we study and discuss the predictability of talk attendance at academic conferences. We start with first statistics concerning the talk and session attendance behavior at the HT 2011 conference.

### 6.2.1 Talk and Session Attendance Statistics

In Table 2 we present first statistics about the talk and session attendance behavior at HT 2011 for the parallel talks. Overall, the 53 conference participants attended 194 sessions. We observe that most of the participants did not change a session during the HT 2011. At this conference, only in 7% ( $\frac{14}{194}$ ) of all cases, the corresponding participant changed the session. In 69% ( $\frac{134}{194}$ ) of all cases the participants visited all talks of the session.


# Sessions	194
# Visited All Talks in Session	134
# Changed Session	14
# Visited exactly 2 talks of Session	13
# Visited exactly 1 talk of Session	33

**Table 2:** Statistics about talk and session attendance behavior at the HT 2011 conference.

### 6.2.2 Influence Factors of Talk Attendance Using Face-To-Face Contact Networks


In this section, we study the influence of face-to-face contacts during a conference on the attendance of talks. Especially we analyze the probability that two participants attended the same talk, based on the current face-to-face contact behavior between these two participants. In the following, we apply a t-test for determining the significance of our observations. We therefore will also plot the 95% confidence intervals of the results. First, we assume that there exists no face-to-face contact between two conference participants until the start of talk  $t$ . In Figure 5, we observe that the probability is nearly random (i. e., probability is 50.8%) that these two participants visit the same talk  $t$ , if there exists no prior face-to-face contact. In addition, we analyze the probability that two participants visit the same talk, when there exists a face-to-face contact till the end of the conference. (Note that this information could not be used for our prediction task, because it used future information.) We observe that the probability here is slightly increased (probability is 55.5%) to attend the same talk, if there will exist a face-to-face contact till the end of the conference. It is interesting to see that the probability is 58.74%, if already a prior face-to-face contact

exists, before the talk starts. This result highlights the influence of face-to-face contacts on the talk attendance. Furthermore, we analyzed whether a face-to-face contact during the coffee break will influence the probability to attend the same talk of the next session. In Figure 5, we see that the probability is 65.5% to attend the same talk of the next session, if there exists a face-to-face contact in the coffee-break before the session.



**Figure 5: Analysis of influence factors concerning the prediction of talks at academic conferences. Here we plot the probability that two participants visit the same talk, given that 1. there is a face-to-face contact in the coffee break before the next talk is going to start, 2. there exists a face-to-face contact before the next talk is going to start 3. there exists a face-to-face contact till the end of the conference, 4. there exists no face-to-face contact.**


In addition, we consider the connection between a conference participant and the presenter of a talk. We study here, whether a participant  $q$  will attend the talk of this presenter  $p$ , when there exists a face-to-face contact between participants  $p$  and  $q$ . In Figure 6, we plot the probability to join the talk of presenter  $p$ , given that there exists a face-to-face contact with presenter  $p$  with minimum contact duration of  $t \geq 20$  seconds (20 seconds is the minimum contact duration). We observe that the probability is 61% that participant  $q$  attends the talk of presenter  $p$ , when there exists a face-to-face contact between participants  $p$  and  $q$ . Note that the probability is 50.8% if there exists no face-to-face contact. When we focus more and more on stronger ties (this means all face-to-face contacts greater than a given time threshold) between these two participants we see that the probability increases almost linearly to attend the talk of presenter  $p$ . Here the probability to attend the talk is 77.78%, if there exists a face-to-face contact with contact duration greater than 960 seconds.



**Figure 6: Given a face-to-face contact between a conference participant  $p$  and a presenter  $q$ . The  $y$ -axis shows the probability that participant  $p$  visits the talk of presenter  $q$ . The  $x$ -axis defines here the minimum face-to-face contact duration between participants  $p$  and  $q$ .**

### 6.2.3 Predictability of Talk Attendance using Simple Baseline Predictors

Next, we analyze the prediction quality of two simple baseline models. We first predict the next talk, based on the number of the accepted papers of the corresponding tracks. This means we predict that a participant joins talk  $A$  (instead of talk  $B$ ), when the number of accepted papers of the track, which talk  $A$  belongs to, is greater than the number of accepted papers for the track, which talk  $B$  belongs to. We see (in Figure 7) that the accuracy of this majority vote predictor is 54.05%. Furthermore, we predict the next talk that the conference participant is going to attend, based on the room of the first talk this participant attended. Here the prediction accuracy is 59.53%.



**Figure 7: Influence factors concerning the prediction of talks at academic conferences. The majority vote predictor predicts the next talk based on greater number of accepted papers for the corresponding tracks.**

### 6.2.4 Predictability of Talk Attendance based on User Interests

We also investigated to what extent conference participants at HT 2011 decided for their attended talks based on the topics of the talks. This is motivated by the general conception that, next to social interactions, personal interest in the presented topics is another major influence factor for talk attendance decisions.

For our analysis, we assume that personal interest is reflected by previous publications. While modeling user interest this introduces limitations with respect to novel upcoming topics, it is based on observable facts and thus leads to simple inference. We downloaded all accessible publications of a user with a publication date before the beginning of the conference. From these, we counted word occurrences into bag-of-word models.

All bag-of-word models were generated after removing stopwords, stemming word tokens using the porter stemmer, and tf-idf weighting. In the most simple setup, we estimate similarities between a visitor's interest and the topic of a talk by calculating the cosine similarity between the respective bag-of-word vectors. For each user and time slot, we predict which out of two parallel talks is attended. The predictor itself is  $\text{argmax}$ , i.e., we predict the talk with the higher cosine similarity to the participants interest model. In order to avoid further influence factors in our experiments, we evaluate content-based influence on a core of our dataset. In this core, only those 51 out of the original 53 users are retained, for which we were able to download at least one prior publication.


To model topics of talks, we build bag-of-word vectors directly from the presented papers in the proceedings. For additional experiments, we also limit to bag-of-word models derived only from abstracts or paper titles.

To find out more about the topical separation of parallel sessions with respect to cosine similarity, we calculated and compared all talk-talk similarities inside and across sessions in the same time-slot.

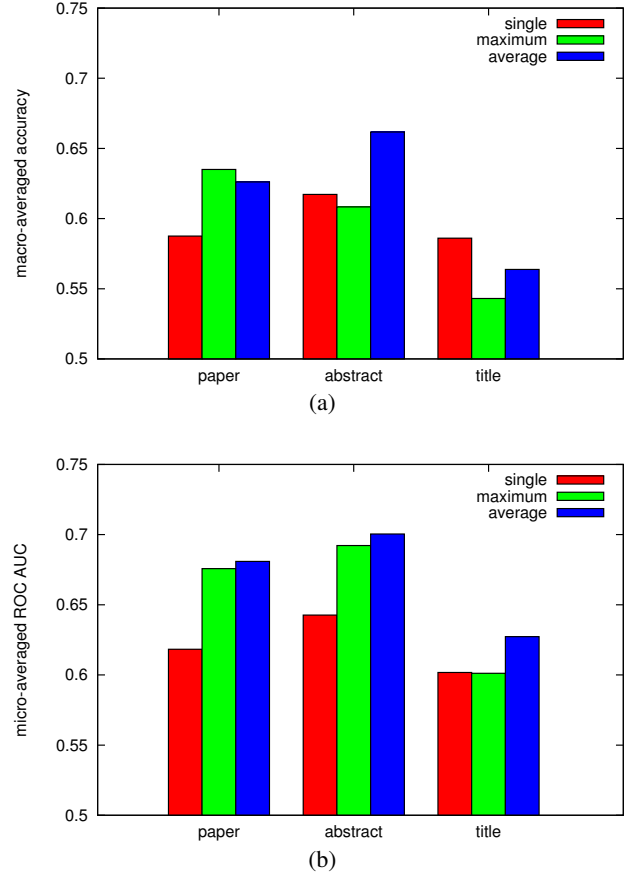
Given our observation that participants at HT 2011 changed between sessions very infrequently, it appears that conference participants decide for attending whole sessions rather than individual talks. If all talks inside the same session have a much stronger topical relation than talks from different sessions, then every individual talk is already a good representative for the session topic. In this case it does not matter which single talk is used to predict the attended session. Otherwise, prediction should be based on all similarities to all the talks in the session. To find out about how well each talk is associated to its session according to the cosine measure, we apply a cluster quality analysis. Here, clustering does not refer to the outcome of a clustering algorithm but to the true distribution of talks over sessions. For each pair of parallel sessions, we calculate the average silhouette value [23] over all talks of both sessions. The silhouette value adopted to our task is defined as  $\text{silh}(t) = \frac{\text{dist}(t, s_{-t}) - \text{dist}(t, s_t)}{\max(\text{dist}(t, s_{-t}), \text{dist}(t, s_t))}$ , where  $t$  is a certain talk,  $s_t$  the session, to which  $t$  belongs, and  $s_{-t}$  is the other session at the same time. We choose single-link cosine distance, i.e.,  $\text{dist}(a, b) = 1 - \cos(a, b)$  where  $\cos$  captures the similarity to the closest non-identical talk within the respective session.

Figure 8 compares the average silhouette values for each set of parallel sessions. We observe, that all values are relatively close

to zero. This means that none of the talk representations gives a good explanation for the distribution of talks over sessions although some talk distributions over session pairs, especially for sessions 11 and 12, show slightly higher silhouette values.



**Figure 8: Cluster quality of the true distribution of talks over parallel sessions; measured by silhouette coefficient calculated using cosine distance based on different talk representations**



**Figure 9: Accuracy (a) and AUC (b) values for content-based talk prediction based on cosine similarity depending on the talk representation and whether the decision is predicted for each talk separately or once for all talks of a session by either maximum or average talk similarity**



Motivated by the intuition that people might make their talk attendance decisions based on one or two talks rather than on all talks of the session, and that participants do not change sessions, we also predict talk attendance session-wise. We tried two options for predicting each users attended session. For the first option, we predict the session with the higher maximum similarity of a talk in the respective session. For the second option, we choose the session with the higher average talk similarity. Figure 9(a) depicts the accuracy values averaged over all 337 individual decisions of a certain participant for a certain talk. Figure 9(b) further shows the area under the ROC curve results. For building up the global micro-average ranking, we first normalized all cosine scores of each particular person and time-slot by dividing by their sum.

As can be seen from Figure 9(b), session-wise prediction constantly achieves superior AUC scores. While this is in accordance to our finding that only few people change between parallel sessions, it is also although the silhouette analysis revealed that the distribution of talks over parallel sessions is not clearly explainable by the cosine-based model. Overall, this can be interpreted as an indication that, from the model’s perspective, the decision of a participant for one of the session is hard, and, that it is usually a consideration of one or more of the most interesting talks. For the global prediction ranking measured by the AUC, the average attendee-talk similarity constantly results in the best results. For the accuracy, the session-wise predictors are also favorable for paper and abstract based talk models. An exception is the title-based model. There both session-wise predictors score below the talk-wise predictor. Title-based models, produce very sparse bag-of-word vectors and might thus lead to low or zero scores for both parallel talks. Yet, the talk-wise predictor achieves a mean accuracy that is only slightly lower compared to other, less sparse models. This indicates that title words are good topic indicators. Lower results for the session-wise predictors, especially the maximum predictor, may be due to wrong decisions based on few overly weighted word matches in one of the talks’ titles. This might also be one reason for the low maximum predictor accuracy on abstract-based talk models. Still, the confidence scores of low quality predictions must be low enough to not influence the global AUC ranking. Interestingly, the results with abstract-based talk models constantly compare favorably to full-paper models in all settings. Potential explanations are the summarizing character of abstracts and the fact that people often do not have the opportunity to read the full paper before choosing the talk to attend. In order to avoid the sparsity problem, we also experimented using dimensionality reduction like for example used by [25] which is similar to *Latent Semantic Indexing* [10]. SVD is used here to map to a denser lower dimensional vector space spanned by the  $k$ th largest eigenvalues of a matrix  $MM^T$ . The columns of  $M$  are the sums of all context window vectors of a token. Context vectors are derived from sliding windows around each token in the previous papers of the (tracked) participants. However, results did not lead to a clear improvement.

Furthermore, we experimented with a feature-selection based on the pointwise mutual information (PMI) of the two conditional probability estimates for *some author used term  $t$  in his previous papers* and *some author used term  $t$  in the proceedings* given a term  $t$ . This PMI was intended to measure cosine similarity only based on tokens which are to some extend related to the topics of the conference. However, this also did not yield a better interest model for explaining talk attendances. A potential cause is the small number of presentations and the resulting low amount of data for estimating the probabilities.

### 6.2.5 Predictability of Talk Attendance Using the Hybrid Rooted PageRank Predictor

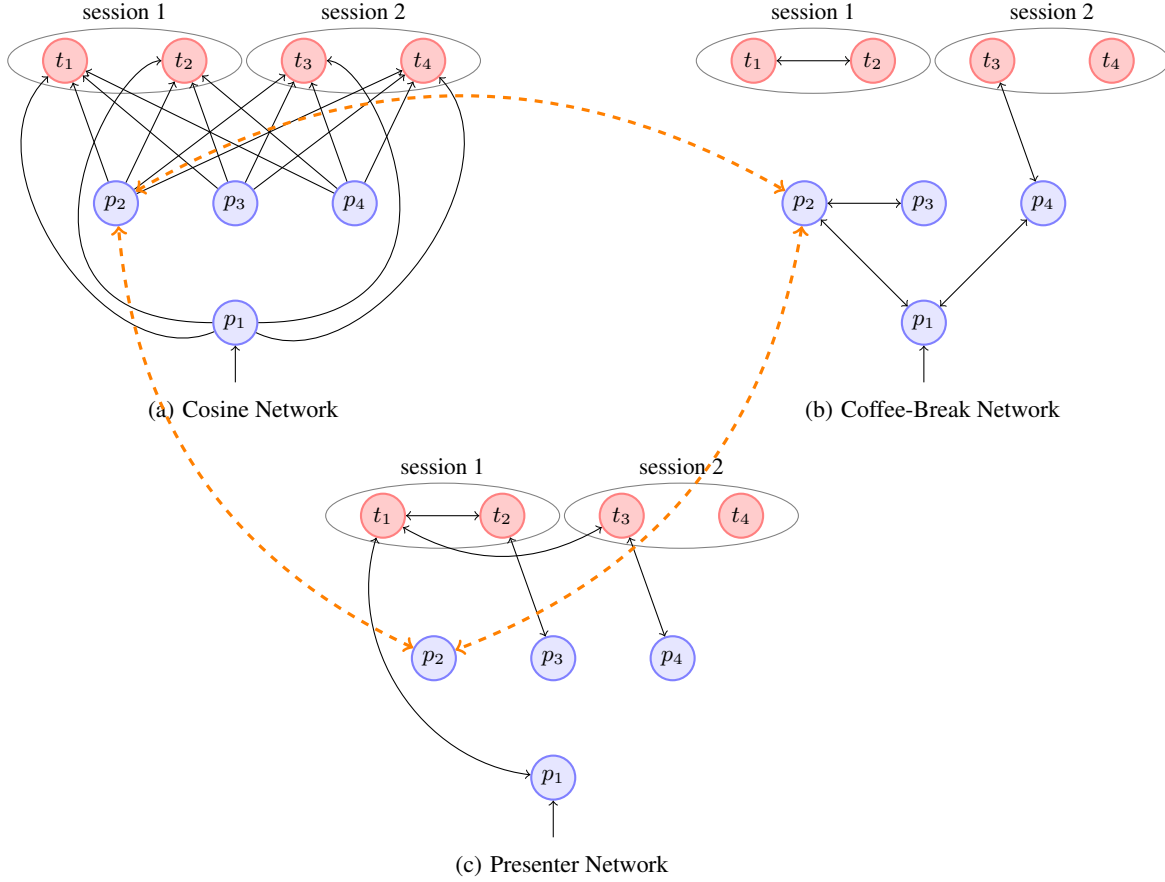
In this subsection, we analyze the predictability of talk attendance using a combination of different networks. For this analysis, we use the *Hybrid Rooted PageRank* (HRPR) algorithm (see section 5.2) as predictor. The advantage of this algorithm is that we can analyze and compare the predictive power of different networks and combinations of these networks. Using the HRPR-algorithm, we combine the information of the paper-similarity network, the aggregated face-to-face contact network (of the coffee break before the next talk is going to start), and the presenters face-to-face contact network. The structure of these graph is illustrated in Figure 10. Note that the hybrid rooted random walk (see Algorithm 1) selects a network with respect to a given probability distribution  $P = (p_1, p_2, p_3)$ . In our experiments, we studied all parameter combinations with  $p_1 + p_2 + p_3 = 1$  and  $p_1, p_2, p_3 \in \{0.0, 0.1, 0.2, \dots, 1.0\}$ . Assume we want to predict, whether participant  $p$  attends talk  $t_1$  or talk  $t_2$ . The predicted talk is then given by the talk  $t_i$ , where  $p_i$  is the presenter of talk  $t_i$ ,  $p_j$  the presenter of talk  $t_j$  and  $HRPR(p, p_i) > HRPR(p, p_j)$ .

We start by analyzing the predictive power for each network separately. In Figure 11 a) and b), we observe that the paper-similarity network performs best with an AUC-value of 0.630 and an accuracy of 0.610. These results correspond to the results of the single variant in Figure 9). Using just the face-to-face contact network of the coffee-break does not work as well as using the paper-similarity network. Here the AUC-value is 0.596. In contrast to the coffee break’s face-to-face contact network, the presenter network contains just the links from the presenter of the next talk that is going to start. We observe that using just the presenters face-to-face contact network does not perform very well and works worse than using just the face-to-face contact network. This is because most participants do not have a face-to-face contact to a presenter before the presenter’s talk starts. Hence, the presenter network is rather sparse and does not provide major predictive power on its own. The AUC-value for the presenter network is 0.474. In this context, the observation that the face-to-face contact network works better than the presenter networks suggests that links between participants help further to improve the predictive power.

Furthermore we analyze, whether the combination of different networks increases the predictability of talk attendance at academic conferences. In Figure 11, we observe that the best result can be obtained by combining the information of all networks. However, the increase of predictability by combining the information of different networks is rather small, and we do not know the parameter combinations leading to the best results. The result just gives an indication that a combination can help to increase prediction quality. In our analysis, we handle the presenter network as an additional network. This gives us the possibility to weight a link between a participant and presenter separately. We observe here that the predictability could not be increased when we combine the presenter network and the face-to-face contact network.

In Table 2, we observed that most participants visited all talks in one session. Furthermore, it was unlikely that a participant changed a session. Despite this observation, it is natural to assume that a participant is not interested in each talk of one session. We argue here that, in most cases, at most one or two talks of a session are the cause for attending the session. Therefore, for each network, we merge the nodes of all presenters in one session. The merged nodes thus represent the whole session. The weight vectors for in- and





**Figure 10: Example illustrations of the *Hybrid Rooted PageRank* networks for talk prediction.** Graph (a) shows the structure of a cosine (user interests) network connecting persons  $p_1, \dots, p_4$  with the presenters of the talks  $t_1, \dots, t_4$  for which attendance is to be predicted. Graph (b) shows the coffee-break network linking persons for whom face-to-face contact have been measured in the coffee-break before the talk. Graph (c) is the presenter network, containing face-to-face contacts at any time before the talks to be predicted. Dashed links show an example part for switching the networks. Such links exist between all nodes with equal labels.

out-going edges are calculated as the re-normalized sum of the respective individual nodes' weight vectors. The merged network results depicted in Figure 12 clearly show an increase in talk prediction quality. Considering the best tested parameter combinations, the AUC score increases from 0.638 to 0.703 and accuracy increases from 0.617 to 0.666. We also observe that, for each parameter combination, the combination of all networks performs better, when we merge the presenter nodes. Unlike the model where we do not merge the presenter nodes of one session, we observe that the combination of the presenter network and the face-to-face contact network increases the prediction accuracy significantly, when we merge the presenter nodes. Considering the best parameter combinations for the presenter and face-to-face contact network, the prediction quality increases from 0.611 to 0.68 AUC. A further interesting point is that a minimal fraction of the face-to-face contact network or paper-similarity network increases the predictive power of the presenters face-to-face contact network from 0.61 to 0.661 and 0.655. For our surprise, this trend can not be observed for the face-to-face contact network results. In addition, we observe that our presented approaches significantly outperform the simple predictors presented in Section 6.2.3.

## 7. CONCLUSIONS

In this paper, we analyzed and discussed the predictability of talk attendance at academic conferences evaluated on real-world data. We considered different influence factors, concerning this prediction problem. Specifically, we studied the influence of face-to-face contacts and user interests on the talk attendance. We showed that the probability of two participants attending the same talk is nearly random, if there exists no face-to-face contact before the talk is going to start. In this context, the probability (that two participants attend the same talk) is significantly increased if there exists a face-to-face contact in the break before the talk. Next, we analyzed the influence of user interest on talk attendance. We observed, that prediction based on user-interest alone achieves better results than prediction based solely on face-to-face contact data. Using the *Hybrid rooted PageRank* we showed that a combination of different networks helps to further improve the prediction accuracy. Another important observation is that the combination of all information belonging to one session, i. e., merging the presenter nodes, significantly improves prediction accuracy. This supports the theory that, in many cases, only one or two talks are the cause for a user's session attendance decision. Since it is unknown which of the talks in a session is relevant for the decision, all have to be considered.

		AUC values, HRPR											
		PaperSim	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Presenter													
0			59.6	58.2	60.4	61.1	62.8	62.8	63.0	62.6	63.1	63.3	63.0
0.1			57.3	61.9	62.4	62.7	63.5	63.6	63.1	63.6	63.1	63.6	
0.2			57.3	63.1	62.7	63.1	63.2	63.0	63.5	63.8	63.6		
0.3			56.6	63.8	63.4	63.7	62.5	63.0	62.9	63.7			
0.4			56.7	63.3	63.1	63.0	62.4	62.9	63.1				
0.5			56.1	63.4	62.5	61.9	61.9	61.2					
0.6			55.5	63.1	62.0	60.8	61.5						
0.7			56.1	61.4	60.7	60.5							
0.8			54.6	59.8	58.8								
0.9			52.7	56.0									
1			47.4										

(a) Single Variant

		ACC values, HRPR											
		PaperSim	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Presenter													
0			55.6	56.3	59.1	60.0	61.6	61.1	60.9	60.6	60.2	60.2	61.0
0.1			55.6	59.2	58.2	60.7	61.1	61.3	60.4	59.7	59.1	58.8	
0.2			55.7	59.5	60.7	60.3	59.7	59.5	57.9	58.4	58.8		
0.3			56.8	61.3	61.0	59.7	58.4	57.7	57.7	58.9			
0.4			56.5	61.3	59.7	59.5	58.2	58.1	58.9				
0.5			55.7	61.7	59.6	58.8	57.4	58.1					
0.6			54.2	60.9	58.6	56.5	58.1						
0.7			54.6	58.1	56.8	57.4							
0.8			53.1	57.2	56.7								
0.9			52.2	54.3									
1			48.6										

(b) Single Variant

Figure 11: AUC and accuracy values for the talk prediction task using the *Hybrid Rooted PageRank* as predictor. The  $x$ -axis represents the probability to choose the paper-similarity network in the random walk of *HRPR*, the  $y$ -axis the probability to choose the presenter network. The probability to choose the coffee-break network is then defined as  $1-x-y$ . The  $z$ -axis displays the AUC-value for the defined (by the  $x$  and  $y$  axes) parameter combinations. In these figures we present the predictability-results, without merging the presenter nodes.

		AUC values, HRPR											
		PaperSim	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Presenter													
0			61.1	56.6	57.4	60.2	62.4	64.7	66.2	68.0	68.5	68.1	67.0
0.1			68.5	62.5	63.4	65.3	67.4	68.9	69.3	68.8	68.8	70.3	
0.2			67.1	66.5	66.5	68.1	68.5	68.9	69.5	69.3	69.2		
0.3			65.5	67.9	68.2	68.7	68.7	68.2	68.0	69.2			
0.4			67.8	69.0	69.4	68.8	68.5	67.9	67.5				
0.5			67.0	69.5	69.0	68.4	67.2	67.4					
0.6			68.2	69.8	68.9	67.8	66.4						
0.7			67.5	69.3	68.4	65.7							
0.8			66.8	68.9	66.2								
0.9			66.1	65.5									
1			61.0										

(a) Merged Variant

		ACC values, HRPR											
		PaperSim	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Presenter													
0			57.5	56.5	54.3	58.2	60.2	61.0	60.4	62.8	62.0	61.0	64.2
0.1			62.1	58.5	59.6	62.5	64.5	63.4	64.1	63.5	62.8	65.7	
0.2			61.7	61.0	61.4	63.6	64.2	64.5	63.1	61.7	64.2		
0.3			60.4	63.4	63.0	63.2	64.8	62.8	61.3	63.8			
0.4			62.5	63.6	64.1	65.9	66.0	62.7	63.4				
0.5			61.7	64.3	65.5	65.6	63.2	63.8					
0.6			62.1	66.6	65.6	64.1	63.5						
0.7			61.7	65.3	66.0	62.5							
0.8			62.3	66.6	63.2								
0.9			62.1	62.7									
1			57.7										

(b) Merged Variant

Figure 12: AUC and accuracy values for the talk prediction task using the *Hybrid Rooted PageRank* as predictor. The  $x$ -axis represents the probability to choose the paper-similarity network in the random walk of *HRPR*, the  $y$ -axis the probability to choose the presenter network. The probability to choose the coffee-break network is then defined as  $1-x-y$ . The  $z$ -axis displays the AUC/Accuracy-value for the defined (by the  $x$  and  $y$  axes) parameter combinations. In these figures we present the predictability-results, when we merge the presenter nodes.

## 8. REFERENCES

- [1] G. Adomavicius and A. Tuzhilin. Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *Knowledge and Data Engineering*, 17(6), 2005.
- [2] H. Alani, M. Szomszor, C. Cattuto, W. V. den Broeck, G. Correndo, and A. Barrat. Live Social Semantics. In *Intl. Semantic Web Conference*, pages 698–714, 2009.
- [3] M. Atzmueller, M. Becker, M. Kibanov, C. Scholz, S. Doerfel, A. Hotho, B.-E. Macek, F. Mitzlaff, J. Mueller, and G. Stumme. Ubicon and its Applications for Ubiquitous Social Computing. *New Review of Hypermedia and Multimedia*, 20(1):53–77, 2014.
- [4] M. Atzmueller, S. Doerfel, A. Hotho, F. Mitzlaff, and G. Stumme. Face-to-Face Contacts at a Conference: Dynamics of Communities and Roles. In *Modeling and Mining Ubiquitous Social Media*, volume 7472 of *LNAI*. Springer Verlag, Heidelberg, Germany, 2012.
- [5] A. Barrat, C. Cattuto, V. Colizza, J.-F. Pinton, W. V. den Broeck, and A. Vespignani. High Resolution Dynamical Mapping of Social Interactions with Active RFID. *CoRR*, abs/0811.4170, 2008.
- [6] A. Barrat, C. Cattuto, M. Szomszor, W. V. den Broeck, and H. Alani. Social Dynamics in Conferences: Analyses of Data from the Live Social Semantics Application. In *International Semantic Web Conference (2)*, pages 17–33, 2010.
- [7] S. Brin and L. Page. The Anatomy of a Large-Scale Hypertextual Web Search Engine. *Computer Networks*, 30(1-7):107–117, 1998.
- [8] C. Cattuto, W. Van den Broeck, A. Barrat, V. Colizza, J.-F. Pinton, and A. Vespignani. Dynamics of Person-to-Person Interactions from Distributed RFID Sensor Networks. *PLoS ONE*, 5(7):e11596, 07 2010.
- [9] Christoph Scholz and Martin Atzmueller and Alain Barrat and Ciro Cattuto and Gerd Stumme. New Insights and Methods For Predicting Face-To-Face Contacts. In *Proc. 7th Intl. AAAI Conference on Weblogs and Social Media*, 2013.
- [10] S. Deerwester, S. Dumais, G. Furnas, T. Landauer, and R. Harshman. Indexing by Latent Semantic Analysis. *Journal of the American Society for Information Science* 41, pages 391–407, 1990.
- [11] N. Eagle, A. Pentland, and D. Lazer. From the Cover: Inferring Friendship Network Structure by using Mobile Phone Data. *Proceedings of The National Academy of Sciences*, 106:15274–15278, 2009.
- [12] J. A. Hanley and B. J. McNeil. The Meaning and Use of the Area under a Receiver Operating Characteristic (ROC) Curve. *Radiology*, 143(1):29–36, Apr. 1982.
- [13] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot. Pocket Switched Networks and Human Mobility in Conference Environments. In *Proceedings of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*, WDTN '05, pages 244–251, New York, NY, USA, 2005. ACM.
- [14] L. Isella, M. Romano, A. Barrat, C. Cattuto, V. Colizza, W. Van den Broeck, F. Gesualdo, E. Pandolfi, L. Ravà, C. Rizzo, and A. Tozzi. Close Encounters in a Pediatric Ward: Measuring Face-to-Face Proximity and Mixing Patterns with Wearable Sensors. *PLoS ONE*, 6:e17144, 2011.
- [15] L. Isella, J. Stehlé, A. Barrat, C. Cattuto, J.-F. Pinton, and W. V. D. Broeck. What's in a Crowd? Analysis of Face-to-Face Behavioral Networks. *Journal of Theoretical Biology*, 271:166–180, 2011.
- [16] T. Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '02, pages 133–142, New York, NY, USA, 2002. ACM.
- [17] D. Lee and P. Brusilovsky. Exploring Social Approach to Recommend Talks at Research Conferences. In *COLLABORATECOM 2012 - 8th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing*, Oct. 2012.
- [18] D. Liben-Nowell and J. M. Kleinberg. The Link Prediction Problem for Social Networks. In *CIKM*, pages 556–559, 2003.
- [19] B.-E. Macek, C. Scholz, M. Atzmueller, and G. Stumme. Anatomy of a Conference. In *Proc. 23rd ACM Conference on Hypertext and Social Media*, pages 245–254, New York, NY, USA, 2012. ACM Press.
- [20] M. Meriac, A. Fiedler, A. Hohendorf, J. Reinhardt, M. Starostik, and J. Mohnke. Localization Techniques for a Mobile Museum Information System. In *Proceedings of WCI*, 2007.
- [21] E. Minkov, B. Charrow, J. Ledlie, S. Teller, and T. Jaakkola. Collaborative Future Event Recommendation. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*, CIKM '10, pages 819–828, New York, NY, USA, 2010. ACM.
- [22] M. C. Pham, D. Kovachev, Y. Cao, G. M. Mbogob, and R. Klamma. Enhancing Academic Event Participation with Context-aware and Social Recommendations. In *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*, ASONAM '12, pages 464–471, Washington, DC, USA, 2012. IEEE Computer Society.
- [23] P. Rousseeuw. Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis. *Journal of Computational and Appl. Mathematics*, 20(1):53–65, 1987.
- [24] C. Scholz, S. Doerfel, M. Atzmueller, A. Hotho, and G. Stumme. Resource-Aware On-Line RFID Localization Using Proximity Data. In *Proc. ECML/PKDD 2011*, 2011.
- [25] H. Schütze. Automatic Word Sense Discrimination. *Computational Linguistics*, 24(1):97–123, 1998.
- [26] J. Stehlé, N. Voirin, A. Barrat, C. Cattuto, L. Isella, J.-F. Pinton, M. Quaggiotto, W. Van den Broeck, C. Régis, B. Lina, and P. Vanhems. High-Resolution Measurements of Face-to-Face Contact Patterns in a Primary School. *PLoS ONE*, 6(8):e23176, 08 2011.
- [27] C. van Rijsbergen, S. Robertson, and M. Porter. New Models in Probabilistic Information Retrieval. 1980.
- [28] C. Wongchokprasitti, P. Brusilovsky, and D. Para. Conference Navigator 2.0: Community-Based Recommendation for Academic Conferences. In *Proc. Workshop Social Recommender Systems, IUI'10*, 2010.
- [29] B. Xu, A. Chin, H. Wang, L. Chang, K. Zhang, F. Yin, H. Wang, and L. Zhang. Physical Proximity and Online User Behavior in an Indoor Mobile Social Networking Application. In *Proc. 4th IEEE Intl. Conf. on Cyber, Physical and Social Computing (CPSCoM 2011)*, 2011.
- [30] X. Zuo, A. Chin, X. Fan, B. Xu, D. Hong, Y. Wang, and X. Wang. Connecting People at a Conference: A Study of Influence Between Offline and Online Using a Mobile Social Application. In *Connecting People at a Conference*, 2012.